

UNITED STATES PATENT APPLICATION FOR

SYSTEM AND METHOD FOR MONITORING PROCESSES OF AN INFORMATION  
TECHNOLOGY SYSTEM

Inventor:

Jerry Champlin

EXPRESS MAIL CERTIFICATE OF MAILING

"Express Mail" mailing label number EV 32880826 US

5 Date of Deposit April 7, 2004

I hereby certify that this paper or fee is being deposited with the United States  
Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on  
the date indicated above and is addressed to the Assistant Commissioner for  
10 Patents, Washington, D.C. 20231.

Norma E. Gillespie  
(Typed or printed name of person mailing paper or fee)

*Norma E. Gillespie*  
(Signature of person mailing paper or fee)

**SYSTEM METHOD FOR MONITORING PROCESSES OF AN INFORMATION  
TECHNOLOGY SYSTEM**

Inventor:  
Jerry Champlin

COPYRIGHT NOTICE

A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention generally relates to monitoring computer systems and, more particularly, to comprehensive and user-friendly monitoring tools for system managers of information technology (IT) systems.

Discussion of Background

Information technology (IT) systems need monitoring in order for the IT systems to work properly. The behavior of IT systems also needs analysis in order to predict any future problems or failures. Monitor systems typically display status information of an IT system on a web page for example. Unfortunately, monitor systems have lacked a comprehensive

user-friendly framework that allows system managers to easily detect and predict current and potential system problems.

## SUMMARY OF THE INVENTION

It has been recognized that what is needed is a monitor system that allows system managers to easily detect and predict current and potential system problems. Broadly speaking, the present invention fills these needs by providing a comprehensive system and method for monitoring processes of an information technology (IT) system. It should be appreciated that the present invention can be implemented in numerous ways, including as a process, an apparatus, a system, a device or a method. Several inventive embodiments of the present invention are described below.

In one embodiment, a system for monitoring processes of an information technology (IT) system is provided. The system comprises a monitor agent configured to collect performance and availability metrics associated with at least one of a host machine, a network, an operating system, a database, and an application; a data loader, wherein the monitor agent is further configured to transmit the metrics to the data loader; an escalation server configured to receive and manage alerts generated by the monitor agent, and further configured to group an alert entering the escalation server into a resource group; and an analysis tool including an analysis tool application configured to assist a system manager in visualizing and understanding the performance of the information technology system through the use of at least one of a visual graph, a performance report, a real-time operating

status, and a system health report. A document center is provided that captures, in a central repository, performance reports, system health reports and any other documentation required by the user. Key performance indicators (KPI) is  
5 provided to rollup data from multiple hosts to provide a summary analysis of performance across all of those hosts.

In another embodiment, a method of monitoring processes of an information technology (IT) system is provided. The method comprises collecting via a monitor agent performance  
10 and availability metrics associated with at least one of a host machine, a network, an operating system, a database, and an application; transmitting the metrics from the monitor agent to a data loader; transmitting alerts from the monitor agent to an escalation server, wherein the escalation server  
15 is configured to group an alert entering the escalation server into a resource group; and analyzing the metrics and alerts using an analysis tool that includes an analysis tool application configured to assist a system manager in visualizing and understanding the performance of the  
20 information technology system through the use of at least one of a visual graph, a performance report, a real-time operating status, and a system health report.

The invention encompasses other embodiments of a system, a method, an apparatus, and a computer-readable medium, which  
25 are configured as set forth above and with other features and alternatives.

## BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be readily understood by the following detailed description in conjunction with the accompanying drawings. To facilitate this description, like  
5 reference numerals designate like structural elements.

FIG. 1 is a schematic diagram of the component architecture of the system, in accordance with an embodiment of the present invention;

FIG. 2 is a flowchart of the activities of the monitor  
10 agent, in accordance with an embodiment of the present invention;

FIG. 3 is an example of a specification for an extensible markup language (XML) document type definition (DTD), in accordance with an embodiment of the present invention;

15 FIG. 4 is a simplified core data model describing the key elements of the data loader, in accordance with an embodiment of the present invention;

FIG. 5 is a simplified class diagram of the composition entities of the data loader, in accordance with an embodiment  
20 of the present invention;

FIG. 6 is a schematic diagram showing the relationships between alert escalation entities, in accordance with an embodiment of the present invention;

FIG. 7 is a flowchart of the escalation management  
25 process, in accordance with an embodiment of the present invention;

FIG. 8 is a flowchart of the report generation process,  
in accordance with an embodiment of the present invention;

FIG. 9 is an example page from a sample report generated  
from the process of FIG. 8, in accordance with an embodiment  
5 of the present invention;

FIG. 10 is an example of a portal system summary screen,  
in accordance with an embodiment of the present invention;

FIG. 11 is an example of a hierarchy view of the portal  
system, in accordance with an embodiment of the present  
10 invention;

FIG. 12 is an example of a graph that contains  
information multiple metrics and multiple hosts, in accordance  
with an embodiment of the present invention;

FIG. 13 is an example of an escalation management  
15 interface, in accordance with an embodiment of the present  
invention;

FIG. 14 is an example of an escalation management  
interface having an escalation list view, in accordance with  
an embodiment of the present invention;

20 FIG. 15 is an example of an escalation management  
interface having an escalation detail view, in accordance with  
an embodiment of the present invention;

FIG. 16 is a simplified Entity Relationship Diagram (ERD)  
showing the general relationship of key performance indicator  
25 metrics, in accordance with an embodiment of the present  
invention;

FIG. 17 shows examples of a web-based interface used for populating a service group hierarchy and associated data, in accordance with an embodiment of the present invention; and

FIG. 18 is an example of the display page for key  
5 performance indicators, in accordance with an embodiment of the present invention.



## DESCRIPTION OF THE PREFERRED EMBODIMENTS

An invention for a system and method for monitoring processes of an information technology (IT) system is disclosed. Numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be understood, however, to one skilled in the art, that the present invention may be practiced without some or without all of these specific details.

### 10 General Overview

The computer system of the present invention is a comprehensive software framework, which provides monitoring, analysis, and management capabilities for client servers and applications through a multi-tier architecture. The various components of this system are designed to be secure, highly available, fault-tolerant, extensible and scaleable.

FIG. 1 is a schematic diagram of the component architecture 101 of the system, in accordance with an embodiment of the present invention. The individual components of the system may be aggregated by the service role they play in the overall framework.

The monitoring and measurement capabilities start with the monitor agent 102, a remote agent installed on the host machine. The agent continuously measures the availability and performance of the host operating system, as well as its services and applications. These metrics are in turn

forwarded to the data loader 104, where they are then processed and made available for the framework analysis applications. The monitor agent 102 may be configured to run on any industry accepted or widely used operating system.

5 Accordingly, the system is platform agnostic.

In addition, the monitor may be configured with thresholds for certain metrics, which, when exceeded, trigger alerts that are sent to the alert escalation server 106. The alert escalation server 106 utilizes a highly configurable set  
10 of rules that determine the notification frequency, escalation path and recipients of each received alert.

The software framework also provides a robust set of tools to analyze as well as manage the large amount of raw data generated by the monitor. The data analysis tools 109  
15 serve to aggregate and condense the data for use in a variety of analysis formats. These tools include automatically generated performance analysis reports 110, key performance indicators (KPI) 112, on-demand trend graphing capability, and real-time status reports on the health of the client system.

20 Another set of tools serve to manage the activity of the monitoring system, which are grouped as the management tools 113. Generally speaking, these tools enable the administration of the configuration of the alert escalation system as well as interacting with the operation of the alert  
25 system itself, allowing users to search, close, suspend and acknowledge escalations generated by the monitor agents.

## The Monitor Agent

The primary function of the monitor agent 102 is to collect performance and availability metrics on the host machine and report them to the data loader 104. In the event  
5 that the monitor agent 102 encounters a measurement or a trend in measurements that exceeds a configured performance threshold or performance trend rule, the monitor agent 102 is also able to generate an alert, which is sent to and handled by the alert escalation server 106.

10 The monitor agent 102 is run as a daemon process and loops through a list of metrics to collect data, as dictated by a time interval specified in the monitor agent's configuration file. The monitor agent 102 itself is designed to be a generic monitoring tool that provides a set of  
15 facilities or application program interface (API) for reporting metrics and handling alerts. However, the metrics themselves are collected by a set of specialized monitor classes, which are loaded, initialized, and executed by the monitor agent 102 during run-time. In this manner, the agent  
20 may be extended to collect additional metrics with little impact on the existing code. The monitor agent 102 also monitors text based log files and generates alerts based on pattern matches or pattern match frequencies which exceed a configured performance threshold.

25 At specified intervals, the monitor agent 102 executes the monitor method of each configured monitoring class, aggregating the complete set of measurements to report back to

the data loader server 104. This set of data is serialized into an extensible markup language (XML) stream for transport via either hypertext transfer protocol (HTTP or HTTPS) or simple mail transfer protocol (SMTP). SMTP is usually  
5 configured as a backup protocol to HTTP or HTTPS for fault-tolerance. In the event that neither protocol succeeds, messages are spooled by the agent until a connection can be re-established. At that time the backlog is gradually processed until clear.

10 As the agent processes each set of metrics, the metric values are compared to the configured alert thresholds. Two distinct thresholds may be set, one for a "warning" condition and another for a "critical" condition. Should either of these thresholds be exceeded, the agent will construct an  
15 alert message, serialized in an XML stream, and send it to the Alert Escalation Server via HTTP or HTTPS. Should the alert fail to be received for whatever reason, the alert will then be transmitted through SMTP as part of a failsafe notification mechanism.

20 FIG. 2 is a flowchart of the activities of the monitor agent 102 described above, in accordance with an embodiment of the present invention.

The data transport relationship established between the monitor agent 102 and the data loader 104 is based on the  
25 design pattern idea of Proxy/Adapter pairs, where a data loader API proxy exists for each transport protocol. The protocol adapter on the data loader 104 is responsible for

deserializing the XML message sent by the proxy and executing the requested operation. This enables a flexible and extensible transport mechanism for communication.

FIG. 3 is an example of a specification for an extensible markup language (XML) document type definition (DTD), in accordance with an embodiment of the present invention.

#### The Data Loader

The data loader 104 is responsible for receiving metric raw data reports transmitted by the remote monitor agents by way of either HTTP or HTTPS or SMTP. For HTTP messages, the XML stream is received and deserialized by a perl module written for the Apache mod\_perl environment. As each metric is received, it is stored in the database and related with the same metrics collected earlier from that same host. An in-memory caching system is used to lookup these metric-host groupings while minimizing database traffic.

Metrics are described within the data repository 108 according to a hierarchical metric taxonomy, which conceptually relates classes of metrics with one another. For example, all data related to the host operating system are differentiated from data related to hosted applications. Furthermore, each of these branches is further refined and classified into sections - Disk activity and CPU activity within the operating system branch for example. These metric paths are in turn associated with individual hosts for which corresponding data is collected. In this way, collected data

can be cataloged to a particular path and host and retrieved for subsequent analysis.

The individual metric data paths that describe the data gathered for a host are also tied to the notion of a generic data path, which are not tied to any host but rather describe a general family of data paths. An example of such a data path might be all measurements related to the Apache web server or perhaps all metrics related to disk swapping activity.

FIG. 4 is a simplified core data model describing the key elements of the data loader 104, in accordance with an embodiment of the present invention.

This process is how the raw data for monitored hosts gets stored in the data repository 108. In addition, status information from the latest metric received is stored as a means to display real-time system health information through the data analysis tools 109.

As mentioned above, the HTTP Data Loader has a parallel component in the form of a stand-alone daemon which continuously scans a mail spool for incoming messages via SMTP. The operation of this daemon is, in all other respects, the same as the HTTP loader. The two symmetric processes correspond to the protocol adapters for each proxy/adaptor pair in the system.

FIG. 5 is a simplified class diagram of the composition entities of the data loader 104, in accordance with an embodiment of the present invention. FIG. 5 shows the

relationship of the two symmetric processes of the HTTP proxy and the SMTP proxy.

The HTTP Loader is also extremely fault-tolerant. In the event the process is interrupted or an exception is

5 encountered during the processing of a message, the message is redirected to the SMTP spool for deferred processing. In the event of a performance degradation, the HTTP loader will also run in an "economy" mode, which defers message processing to the SMTP spool for resource conservation until normal  
10 operating conditions resume.

Because SMTP processing is by nature asynchronous, the existence of this fallback processing mechanism ensures a minimum level of availability given potentially fluctuating system resources.

15

#### Alert Escalation Server

The alert escalation server 106 is the framework system responsible for receiving and managing the alerts generated by the monitor agents 102. Alerts entering the framework are  
20 grouped together by the type of resource generating the alert. These groupings are called resource groups. Examples of resource groups may include database related alerts, application related alerts or operating system alerts. These resource groupings, when assigned to a host and a list of  
25 alert recipients, allow the creation of escalation paths, which represent the lifecycle of an alert incident for a monitored host.

The escalation paths are defined by a sequence of path steps, which progress the alert through its lifecycle. At each sequential step in the path, if the alert is not resolved, the escalation will progress to the next step and alert the people responsible for alerts at that step of the cycle. In this manner, alerts can evolve in scope, reach and urgency depending on their duration and origin. Subsequently received alerts, if originating from the same host for the same resource group are grouped in with the open escalation since they are related to the first alert.

FIG. 6 is a schematic diagram showing the relationships between alert escalation entities, in accordance with an embodiment of the present invention.

The alert escalation server 106 is governed by two principle processes, the alert adapter and the alert sweeper. The alert adapter is a mod\_perl Apache process, primarily responsible for receiving the serialized XML stream from monitor agents, which signal a problem requiring resolution. Upon receiving an alert, the alert adapter will first check to see if the alert is part of an already open escalation. Should an escalation already be open for the alert's resource group and host, the alert will be bundled with the open escalation and alert notifications will continue to be generated as prescribed by the escalation path steps.

However, if the alert received does not have an already open escalation, a new escalation will be opened on behalf of the alert, which will start a lifecycle of notifications for



this and all subsequent related alerts. In this initial treatment of the alert, an immediate notification is usually sent out to the appropriate parties to indicate that a new alert has been received and that an escalation process has  
5 been started. In the event that the adapter encounters an exception during any of part of this process, the adapter sends an error code back to the sending monitor (in the form of an HTTP response), which describes the nature of the error encountered. The monitor will then failsafe the alert along  
10 with the reason that the initial alert notification attempt failed.

The alert sweeper is responsible for managing the escalations that have been opened by the alert adapter. Managing escalations entails sending out alert notifications  
15 according to the defined intervals at each step as well as advancing escalations to subsequent steps when needed. If an escalation has been manually suspended for any length of time, the sweeper will also see if the suspension duration has expired and the escalation path should be resumed. The alert  
20 sweeper will also automatically close escalations in the event that the resource that generated the alert starts sending in normal, signaling that the problem has been resolved. Management of the escalation will also check to see if the host or a group of hosts in question is being maintained (a  
25 configuration option), which has the effect of suppressing alert notifications as well as escalations.

FIG. 7 is a flowchart of the escalation management process, in accordance with an embodiment of the present invention.

## 5    The Analysis Tools

The analysis tools are a collection of processes, which collaborate to assist managers in visualizing and understanding the performance of their systems through the use of visual graphs, performance reports, real-time operating  
10    status and system health. The function of key performance indicators (KPI) 112 is to rollup data from multiple hosts to provide a summary analysis of performance across all of those hosts. The technologies required to generate these products include the data aggregation process, the custom graphing  
15    engine, the reporting engine and the web portal.

The analysis process starts with the conversion of raw measurement data into aggregated data for various time intervals. Aggregated data records various aspects of the raw data sets for a given duration, including its minimum,  
20    maximum, mean, median, standard deviation, skew, kurtosis and percentile data. This condensed raw data facilitates the manipulation and presentation of measurement data by the tools mentioned above. This process is driven by the rollup daemon, a scaleable, distributable sub-system, which processes  
25    incoming raw data and summarizes it according to time intervals specified by the metric's assigned generic data category. Once this data has been condensed, it is then

available for use by the graphing engine and reporting engine for analysis.

The graphing engine is a collection of perl modules, which provide a programmatic interface to easily map and  
5 manipulate metric data, grouped by data category, into data files and graph definitions. Note the present invention is not limited to the perl scripting language. The programmatic interface may provided by another language, such as C, C++, Java, or any other suitable language.

10 The files can then be used by a custom designed java-based graphing engine to generate sophisticated visualizations of the metric data. By creating definitions for these graphs, managers can see performance trends of their systems as well as establish visual comparative relationships by grouping  
15 related metric data and systems.

The graphing system is used by both the reporting engine as well as the portal for presenting analysis information to system managers.

The reporting engine processes re-usable text templates,  
20 which describe the compositional elements of a performance analysis report. These templates provide a structure for describing not only creating descriptive text for the report, but also for generating various types of graphs with the graphing engine as well as a structured language for

25 heuristically generating data analysis depending on the data's characteristics. This allows the reporting engine to easily generate very detailed and descriptive reports on the

performance of a host's operating system or application as well as provide analysis on the data presented to make recommendations for improving performance or availability.

FIG. 8 is a flowchart of the report generation process, in accordance with an embodiment of the present invention. The system that takes these meta-report templates and constructs them into portable document formatted (PDF) files relies on a number of interrelated technologies. The data aggregation and graphing capability is delegated to the rollup daemon and graphing engine described above. Furthermore, the report templates are parsed using a template processing engine called the Template Toolkit (a perl open-source module). The reports are then assembled into text files formatted for TeX processing, an open source document processing system, which incorporates all the textual and graphical elements into a nicely formatted PDF file.

FIG. 9 is an example page from a sample report generated from the process of FIG. 8, in accordance with an embodiment of the present invention.

The last system component of the suite of analysis tools is the portal, which contains abilities to present real-time and historical system performance information on-demand through a graphical web-based user interface.

The system presents the portal user with a high-level summary of the status or health of the various servers that are currently being monitored. A color-coded system of red, yellow and green quickly alerts her to the overall status and

which group of hosts has any outstanding issues. In order to quickly locate the source of the problem, or just to view the general condition of server metrics that are performing within acceptable thresholds, the user may utilize a "tree-menu" or  
5 collapsible menu, which allows a quick navigation through the hierarchy of metric data being monitored for that host.

FIG. 10 and FIG. 11 are examples of these views that allow a quick navigation through the hierarchy of metric data being monitored for that host. FIG. 10 is an example of a  
10 portal system summary screen, in accordance with an embodiment of the present invention. FIG. 11 is an example of a hierarchy view of the portal system, in accordance with an embodiment of the present invention.

Next to each leaf of the hierarchy is the most current  
15 measurement value for that particular metric. In addition, by selecting that metric, the user is then able to graph the historical data for that metric over any specified length of time, which uses the system graphing engine API. These graphs can further be manipulated to contain multiple metrics (for  
20 comparative analysis) and/or multiple hosts.

An additional feature of the system is a document center that captures, in a central repository, performance reports, system health reports and any other documentation required by the user.

25 FIG. 12 is an example of a graph that contains information multiple metrics and multiple hosts, in accordance with an embodiment of the present invention. Technology

behind this portal user interface includes the Apache™ web server, the mod\_perl extension, the Apache™ PageKit™ web publishing system and custom application business and presentation logic within these frameworks.

5

### Key Performance Indicators

The system framework supplies many ways to view and analyze the raw data collected by the monitor agent. However, up until this point, the analysis tools focused solely on  
10 specific metrics for specific hosts. The ability of the system framework to report on data aggregated by host groups has not yet been discussed. This logical grouping of metrics across hosts may be referred to as "key performance indicators" (KPI) 112. KPI 112 is very useful for performance  
15 analysis as it allows one to quickly measure the performance of overall system application function, availability and health.

For example, a key performance indicator that a user may be interested in tracking is the availability of a web-based  
20 application. Using KPI 112, the user may quickly see the overall health of the application as KPI 112 tracks the system health of all the critical components involved, from the database server, the application servers, the web servers and load balancers. Should any of these components become  
25 unavailable, KPI 112 is capable of inferring that the entire application has been compromised.

In addition, KPI 112 is useful for capacity analysis and planning. Because KPI 112 is capable of aggregating metrics

across groups of hosts, planners can quickly see the amount of disk, CPU and memory utilization and the trends associated with each for their entire hosting environment.

The logical groupings KPI 112 uses to aggregate a set of metrics may be referred to as a "service group". Service groups may contain other service groups but are primarily composed of one or more host data metrics. Should any of the member metrics show a warning or critical status, the overall status of the service group may be affected, which changes the state of the service group's availability. In addition, service groups serve as the basic unit for aggregating core metrics, including CPU, disk, and memory usage. These core metrics may be referred to as "KPI metrics". These KPI metrics are calculated for their related service groups and stored over time according to a specified frequency.

FIG. 16 is a simplified ERD showing the general relationship of KPI metrics, in accordance with an embodiment of the present invention.

The service group hierarchy and associated data is populated using a web-based interface, accessible from the system administrative portal. Using this interface, administrators can create service groups, specify which KPI metrics should be tracked for each group and also create an availability definition composed of host metrics, which will be used in calculating the service group's overall availability.

FIG. 17 shows examples of a web-based interface used for populating a service group hierarchy and associated data, in accordance with an embodiment of the present invention. Once

the service groups and their associated KPI metrics have been created using this interface, the KPI aggregation daemon gathers and computes the results for each service group based on the real-time host metrics sent in by the agent and stored  
5 by the data loader. The KPI aggregation daemon is responsible for calculating KPI metrics for every defined service group according to the frequency specified by the KPI data table. As these values are calculated for each service group per interval, they are stored in the KPI data table where they can  
10 be used for generating KPI analysis graphs in the portal, for example, service group availability for the past 30 days, aggregate CPU utilization for the past 30 days, etc. These analysis graphs may be defined in the KPI administrative area of the system portal, where, once defined by an administrator,  
15 they may be included for display by portal users.

FIG. 18 is an example of the display page for key performance indicators, in accordance with an embodiment of the present invention.

20

#### Management Tools

FIG. 13 is an example of an escalation management interface, in accordance with an embodiment of the present invention.

25

FIG. 14 is an example of an escalation management interface having an escalation list view, in accordance with an embodiment of the present invention.



FIG. 15 is an example of an escalation management interface having an escalation detail view, in accordance with an embodiment of the present invention.

The management tools primarily consist of alert management interfaces available from the portal. The management console allows portal users to interact and manage most levels of the alert escalation server, including the ability view, acknowledge, suspend, or close escalations and their associated alerts, as well as the administrative components of creating and editing escalation paths and their lifecycle. In addition to these escalation and alert management tools, users also have the capability to suppress the alert system altogether by creating and maintaining host maintenance windows, which effectively tell the alert escalation system to ignore alerts generated by that host.

#### System And Method Implementation

Portions of the present invention may be conveniently implemented using a conventional general purpose or a specialized digital computer or microprocessor programmed according to the teachings of the present disclosure, as will be apparent to those skilled in the computer art.

Appropriate software coding can readily be prepared by skilled programmers based on the teachings of the present disclosure, as will be apparent to those skilled in the software art. The invention may also be implemented by the preparation of application specific integrated circuits or by

interconnecting an appropriate network of conventional component circuits, as will be readily apparent to those skilled in the art.

The present invention includes a computer program product  
5 which is a storage medium (media) having instructions stored thereon/in which can be used to control, or cause, a computer to perform any of the processes of the present invention. The storage medium can include, but is not limited to, any type of disk including floppy disks, mini disks (MD's), optical disks,  
10 DVD, CD-ROMs, micro-drive, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, DRAMs, VRAMs, flash memory devices (including flash cards), magnetic or optical cards, nanosystems (including molecular memory ICs), RAID devices, remote data storage/archive/warehousing, or any type of media  
15 or device suitable for storing instructions and/or data.

Stored on any one of the computer readable medium (media), the present invention includes software for controlling both the hardware of the general purpose/specialized computer or microprocessor, and for  
20 enabling the computer or microprocessor to interact with a human user or other mechanism utilizing the results of the present invention. Such software may include, but is not limited to, device drivers, operating systems, and user applications. Ultimately, such computer readable media  
25 further includes software for performing the present invention, as described above.

Included in the programming (software) of the general/specialized computer or microprocessor are software modules for implementing the teachings of the present invention, including, but not limited to, collecting via a  
5 monitor agent performance and availability metrics, transmitting the metrics from the monitor agent to a data loader, transmitting alerts from the monitor agent to an escalation server, and analyzing the metrics and alerts using an analysis tool, according to processes of the present  
10 invention.

In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader  
15 spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.